

NONPARAMETRIC TEXTURE ANALYSIS WITH COMPLEMENTARY SPATIAL OPERATORS

MATTI PIETIKÄINEN AND TIMO OJALA

Machine Vision and Media Processing Unit, Infotech Oulu

P.O. Box 4500, FIN-90014 University of Oulu, Finland

E-mail: {mkip,skidi}@ee.oulu.fi

Recently, we have developed a nonparametric approach to texture analysis, in which the distributions of simple texture measures based on local binary patterns and signed gray level differences are used to provide complementary information about the structural and stochastic properties of image texture. Very good performance has been obtained in various texture classification and segmentation problems. This paper overviews our approach and discusses reasons for its efficiency.

1 Introduction

Texture analysis is important in many applications of computer image analysis for classification, detection, or segmentation of images based on local spatial variations of intensity or color. Many different approaches to texture analysis have been proposed. Among the most widely used texture measures are those derived from gray level cooccurrence matrices or difference histograms. "texture energy" measures obtained by local linear transforms, features based on multi-channel Gabor filtering, wavelets, and Markov random field models [4,6,34,3].

Recently, we have developed a nonparametric approach to texture analysis, in which the distributions of simple texture measures based on local binary patterns and signed gray level differences are used to provide complementary information about the structural and stochastic properties of image texture. Very good performance has been obtained in various texture classification and segmentation problems, see e.g. [7,14-23,25-27,32]. This paper overviews our approach and discusses reasons for its efficiency.

2 Texture Description with Simple Spatial Operators

Consider the 3x3 neighborhood illustrated in Fig. 1, where g_i , $i=0, \dots, 8$, denote the gray values of the respective pixels. We start the derivation of our texture operator by defining texture T as the joint distribution of the gray levels of the nine pixels:

$$T = p(g_0, g_1, g_2, g_3, g_4, g_5, g_6, g_7, g_8) \quad (1)$$

g_4	g_2	g_3
g_5	g_0	g_1
g_6	g_7	g_8

Figure 1. A 3x3 neighborhood, g_i denotes the gray value of pixel i

Without losing information, we can subtract the gray value of the center pixel from the gray values of the surrounding pixels:

$$T = p(g_4 - g_0, g_2 - g_0, \dots, g_8 - g_0) \quad (2)$$

Assuming that the differences $g_i - g_0$ are independent of g_0 , we can factorize Eq. (2):

$$T \sim p(g_0) p(g_4 - g_0, \dots, g_8 - g_0) \quad (3)$$

In practice an exact independence is not warranted, hence the factorized distribution is only an approximation of the joint distribution. However, we are willing to accept the possible small loss in information, as it allows us to achieve invariance with respect to shifts in gray scale. Namely, the distribution $p(g_0)$ in Eq. (3) describes the overall luminance of the image, which is unrelated to local image texture, and consequently does not provide useful information for texture analysis. Hence, much of the information in the original joint gray level distribution about the textural characteristics is conveyed by the joint difference distribution:

$$T \sim p(g_4 - g_0, \dots, g_8 - g_0) \quad (4)$$

2.1 Signed Gray-level Differences

Recently, Ojala et al. [23] showed that an approach based on multidimensional distributions of signed gray level differences of neighboring pixel values is very powerful for texture classification. The advantages of gray-level differences over the traditional co-occurrence method are: (1) the differences fall mainly within a narrower range than the gray levels due to the high correlation between gray levels of adjacent pixels, consequently providing a more compact description of texture:

(2) the signed differences are not affected by changes in mean luminance. In comparison to the commonly used absolute differences, the signed differences contain more information about image texture and consequently are more powerful.

In our experiments, the classification performance of two-, four-, and eight-dimensional difference distributions has been evaluated. By computing co-occurring differences within 3x3-pixel subimages (Fig. 1), we estimate the following distributions,

$$p_2(g_1 - g_0, g_2 - g_0) \quad (5)$$

$$p_4(g_1 - g_0, g_2 - g_0, g_3 - g_0, g_4 - g_0) \quad (6)$$

$$p_8(g_1 - g_0, g_2 - g_0, \dots, g_8 - g_0) \quad (7)$$

The volume of the difference space for an image with G gray levels equals $(2G - 1)^k$, where $k=2,4,8$, corresponding to the distribution we are estimating. If we would straightforwardly describe the difference space with a k -dimensional histogram, we would obtain, even with modest values of G , very large histograms that are computationally expensive and suspect to statistical unreliability. Instead of reducing G , for example, with simple requantization of each co-ordinate, we partition the k -dimensional difference space using vector quantization [20]. For this purpose, we employ a codebook of N k -dimensional codewords, which have indices $n = 0, 1, \dots, N-1$. The codebook is trained with the optimized LVQ1 training algorithm [11], by selecting random vectors from each of the samples in the training set. Our method of vector quantization is computationally much simpler than that used by Valkealahti and Oja [35].

We describe the difference information of a texture sample with a difference histogram. The mapping from the difference space to a difference histogram is straightforward. Given a particular k -dimensional difference vector, the index of the nearest codeword corresponds to the bin index in the difference histogram. In other words, a codebook of N codewords produces a histogram of N bins. The difference histogram of a texture sample is obtained by searching the nearest codeword to each vector present in the sample, and incrementing the bin denoted by the index of this nearest codeword.

2.2 Local Binary Patterns

Signed differences are not affected by changes in mean luminance, hence they are invariant against gray scale shifts. We achieve invariance with respect to the scaling of the gray scale by considering just the signs of the differences instead of their exact values:

$$T = p(s(g_1 - g_0), \dots, s(g_8 - g_0)) \quad (8)$$

where

$$s(x) = \begin{cases} 1, & x \geq 0 \\ 0, & x < 0 \end{cases} \quad (9)$$

If we formulate Eq. (8) slightly differently, we obtain the LBP (Local Binary Pattern) operator introduced by Ojala et al. [19]:

$$LBP = \sum_{i=1}^8 s(g_i - g_0) 2^{i-1} \quad (10)$$

The computation of LBP is illustrated in Fig. 2. The original 3x3 neighborhood is thresholded by the value of the center pixel. The values of the pixels in the thresholded neighborhood are multiplied by the binomial weights given to the corresponding pixels. Finally, the values of the eight pixels are summed to obtain the LBP number for this neighborhood.




Figure 2 Computation of Local Binary Pattern (LBP) and contrast measure C

LBP can be regarded as a simplification of p_8 . It is obvious that LBP contains less textural information than p_8 . The motivation for using LBP instead of p_8 is two-fold: LBP's gray scale invariance and computational simplicity. Effectively, whereas signed differences measure both the spatial organization (pattern) and the contrast (amount) of local image texture, with LBP we intentionally focus only on spatial structure and discard contrast, as it depends on the gray scale.

LBP is by definition invariant against any monotonic transformation of the gray scale, i.e. as long as the order of pixel values stays the same, the output of the LBP operator remains constant. This makes LBP very attractive in situations where the gray scale is subject to changes due to, for example, varying illumination conditions which often have to be coped with in many applications, e.g. in visual inspection. Computational simplicity is another obvious advantage, as there is no need for the quantization of the feature space or other time consuming computations, but the easily calculated 8-bit LBP numbers are simply accumulated into a histogram of 256 bins. This results in a very straightforward and efficient implementation, which may come in handy in time critical applications.

As we noted earlier, LBP does not address the contrast of texture which is important in the discrimination of some textures. For this purpose, we can combine LBP with a simple contrast measure C as shown in Fig. 2 and consider joint occurrences of LBP and C .

He and Wang [8] introduced the Texture Unit operator, which is similar to LBP, but it uses three levels (i.e. two thresholds) instead of two levels used in LBP. This leads to a more inefficient representation and implementation than with LBP (6561 bin values instead of 256), and according to our tests the three-level operator does not perform any better than LBP. The Texture Unit operator usually also needs a user defined delta value for setting the threshold values, which makes it dependent on the gray scale variance.

LBP encodes various simple feature detectors (edge, curve/line, corner, curve/line end, spot/flatness) at different orientations in a single operator. The LBP histogram contains the density of each feature over a given region. An important point is that for each pixel the output of the best-matching feature detector is only utilized. A conventional texture analysis method based on spatial filtering (e.g., Laws' masks, Gabor filters) creates a feature vector for each pixel containing information from both relevant and irrelevant local pattern matches, making this kind of approach less efficient.

Some may find the performance of LBP surprisingly good, given the small support of 3×3 pixels. One may argue that this operator size is by no means adequate, in comparison to the much larger Gabor filter masks that are often used, for example. Actually, the 'built-in' support of the operator is inherently larger than 3×3 pixels, as only a specific limited set of binary patterns can occur next to a particular binary pattern. Further, the histogram of local operator responses incorporates larger scale texture properties. However, if a larger scale analysis is required, it can be accomplished simply by increasing the predicate (i.e. neighborhood size) of the operator. This means that we choose the eight neighbors of the center pixel from the corresponding positions in different neighborhoods (3×3 , 5×5 , 7×7 , etc.) [15].

3 Experiments with Texture Classification

The 32 Brodatz textures used in the experiments are shown in Fig. 3 [2,35]. The images are 256x256 pixels in size and they have 256 gray levels. Each image was divided into 16 disjoint 64x64 samples, which were independently histogram-equalized to remove luminance differences between textures. To make the classification problem more challenging and generic, three additional samples were generated from each sample: a sample rotated by 90 degrees, a 64x64 scaled sample obtained from the 45x45 pixels in the middle of the 'original' sample, and a sample that was both rotated and scaled. Consequently, the classification problem involved a total of 2048 samples, 64 samples in each of the 32 texture categories [35].

The texture classifier was trained by randomly choosing, in each texture class, eight 'original' samples, together with the corresponding 24 transformed samples, as models. The other half of the data, eight 'original' samples and the corresponding 24 transformed samples in each texture class, was used for testing the classifier. In the classification phase a test sample S was assigned to the class of the model M that maximized the log-likelihood measure:

$$L(S, M) = \sum_{n=1}^N S_n \ln M_n \quad (11)$$

where S_n and M_n correspond to the sample and model probabilities of bin n , respectively.

We estimated distributions p_2 , p_4 and p_8 , by partitioning the difference space with a codebook of 384 codewords. The codebook was trained with the standard optimized LVQ1 training algorithm, by selecting 100 random vectors from each of the 1024 samples in the training set.

In experiments, we obtained classification accuracies of 93.3% for p_2 , 95.7% for p_4 and 96.8% for p_8 , respectively [23]. These results are excellent considering the difficulty of the problem. The simple LBP operator achieved an accuracy of 91.2%, i.e. about 5 % less than the much more complex p_8 operator.

For comparison, experiments were carried out with features based on standard GMRF models and Gabor filtering using different mask sizes, with the same implementation as in the MeasTex site [31]. Two different classifiers, the k nearest neighbor classifier and the multivariate Gaussian discriminant, were used for classification. The best result for the GRMF approach was 77.0%, obtained with a standard 6th order symmetric mask and the multivariate Gaussian discriminant classifier. The best combination of GMRF features obtained from features computed using all models from the 1st order to the 7th order achieved 89.3%. The best standard Gabor features, extracted with a filter bank of three different wavelengths and four different orientations in a 11x11 neighborhood and using the Gaussian discriminant classifier, achieved an accuracy of 90.0%. The best

combination of Gabor features obtained from features computed using all mask sizes between 3×3 and 19×19 pixels achieved an accuracy of 91.8%. We were able to improve the accuracy of the Gabor approach from these figures by using both the means and standard deviations of the magnitudes of the filtered images together as texture features, instead of using only the means, as is normally the case. The 'optimized' filter design strategy of Manjunath and Ma [13] also improved the results.

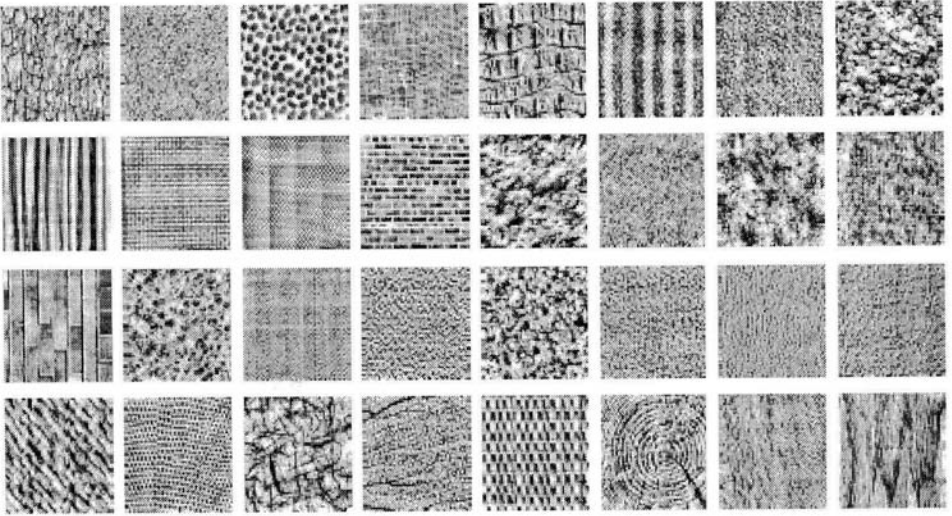


Figure 3. Brodatz textures used in classification experiments.

As we can see, these computationally complex methods used in comparisons performed poorer than our simple methods, even though a larger neighborhood was used to compute the features, instead of the 3×3 neighborhood used in our approach. The rather poor results obtained with the GMRF method are partly caused by the histogram equalization that was used to remove the effects of unequal global brightness and contrast.

4 Experiments with Unsupervised Texture Segmentation

Recently, an unsupervised texture segmentation algorithm using the LBP/C texture measure and nonparametric statistical test was developed by Ojala and Pietikäinen [18]. The method has performed very well in experiments. It is not

sensitive to the selection of parameter values, does not require any prior knowledge about the number of textures or regions in the image as most existing approaches do, and seems to provide significantly better results than existing approaches. The method can be easily generalized, e.g., to utilize other texture features, multiscale information, color features, and combinations of multiple features.

The segmentation method consists of three phases: hierarchical splitting, agglomerative merging and pixelwise classification. First, hierarchical splitting is used to divide the image into regions of roughly uniform texture. Then, an agglomerative merging procedure merges similar adjacent regions until a stopping criterion is met. At this point, we have obtained rough estimates of the different textured regions present in the image, and we complete the analysis by a pixelwise classification to improve the localization. Fig. 4 illustrates the steps of the segmentation algorithm on a 512 x 512 mosaic containing five different Brodatz textures.

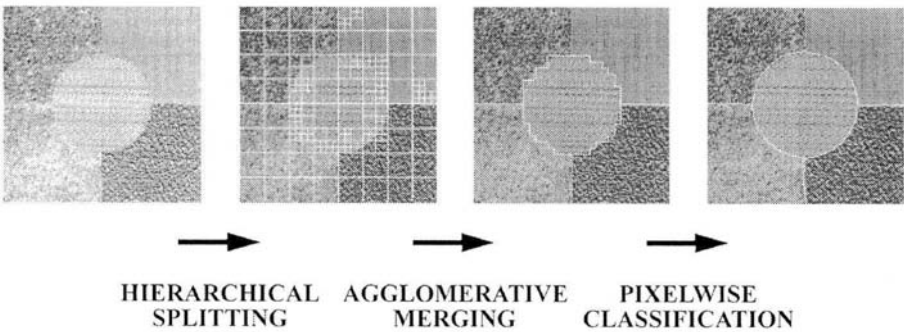


Figure 4. Main sequence of the segmentation algorithm.

We have applied our segmentation method to a variety of texture mosaics with excellent results [18,9]. The database in [9] was created from a set of 86 textures from the Brodatz album, by constructing 100 random mixtures (512x512 pixels each) of five textures. Puzicha et al. [29] detected the overall structure of the mosaic correctly (i.e. the percentage of mislabeled pixels was below 20%) in 95 cases with their ACM (asymmetric clustering model) based multiscale segmentation algorithm when the number of textures was manually provided. The median segmentation error (i.e. the proportion of mislabeled pixels) of all 100 cases was 2.8%, when a topological prior eliminating unlikely label configurations in small neighborhoods was incorporated. Our unsupervised algorithm, which tries to detect the number of textures by itself, found the correct structure in 81 cases, when the minimum region size in hierarchical splitting was set to 32x32 pixels. The median segmentation error of all 100 cases was 2.3%. The 19 failures occurred with mixtures that contained

large-scale texture(s), whose dominating texture pattern was too large to be captured by the 3×3 neighborhood of the LBP:C operator used in our algorithm. This calls for a multiscale approach, where outputs of LBP:C operators computed at different scales are utilized.

The segmentation of a natural scene taken from [24] is shown in Fig. 5 [18]. The textures of natural scenes are generally more non-uniform than the homogeneous textures of the test mosaics. Also, in natural scenes adjacent textured regions are not necessarily separated by well-defined boundaries, but the spatial pattern smoothly changes from one texture to another. Further, we have to observe the infinite scale of texture differences present in natural scenes; choosing the right scale is a very subjective matter. For these reasons there is often no 'correct' segmentation for a natural scene, as is the case with texture mosaics.

The invariance of the LBP:C transform to average luminance shows in the bottom part of the image, where the sea is interpreted as a single region despite the shadows. The result obtained is very satisfactory, considering that important color or gray scale information is not utilized in the segmentation.




Figure 5. Segmentation of a natural scene.

5 Recent developments

In their recent paper, Randen and Husøy reviewed the filtering approaches to texture feature extraction and performed an extensive comparative study using an image set of twelve different types of mosaics of varying complexity [30]. For comparison, the co-occurrence and multiresolution autoregressive (AR) features were also included in the study. In their experiments, various filtering approaches yielded different results for different images. No single approach performed best or very close to the best for all images. The computational complexity of the methods was also considered to be too high. Randen and Husøy concluded that a very useful direction for future research is the development of powerful texture measures that can be extracted and classified with low computational complexity.

We benchmarked the LBP and signed gray level difference approaches in supervised texture segmentation using the same image set as Randen and Husøy and following their experimental setup as closely as possible [23]. The results were excellent. The best signed difference or LBP operator outperformed the best result from Randen and Husøy in 11 of the 12 cases, and in most cases by a clear margin. The LBP approach, despite its simplicity, provided the lowest error rate of all operators in seven of the 12 cases, and the signed gray level difference in four cases, respectively. The impressive result of LBP can largely be attributed to its gray scale invariance. It is understandably a very useful property when the gray scale properties of the unknown test sample differ from the training data, which appeared to be the case in most of the images used in this study.

Recently, we also proposed a simple multiscale extension for LBP by using simultaneously multiple LBP operators computed with different neighborhood sizes [15]. With this approach, the performance of LBP was further improved, indicating that in most cases it is advantageous to use more than just one predicate for the LBP operator. Additionally, we have been developing a multichannel version for color texture analysis, where the computation of LBP is based on an opponent color type representation.

Considering the texture primitive interpretation of LBP as discussed in Section 2.2, it is obvious that all 256 encoded patterns are not equally important. Some patterns should remain more stable under geometric transformations and thus perform better in classification. We have demonstrated that a small subset of local patterns encoded in LBP can perform better than the whole histogram in problems involving geometric transformation between training and testing images [14].

A related observation is that a small subset of 'uniform' local binary patterns with no or a minimal number of spatial $0 \rightarrow 1$ (or $1 \rightarrow 0$) transitions along the circularly symmetric neighborhood of the 8 surrounding pixels appears to sustain rotation very well. Very impressive results were obtained in our recent study on rotation and gray scale invariant texture classification, in which occurrence statistics of these

rotation-invariant ‘uniform’ patterns were used for discrimination [21]. The use of ‘uniform’ patterns means that we ignore ‘nonuniform’ patterns which, by having a larger number of spatial transitions, are more vulnerable to unwanted changes upon rotation. This approximation is supported by the fact that the ‘uniform’ patterns seem to contribute the majority of the local patterns present in Brodatz, like deterministic microtextures, sometimes up to 95%. Further improvement was achieved by quantizing the angular space at a finer resolution (22.5°) than in the original LBP (45°). This kind of operator used in conjunction with a local variance measure outperformed the existing rotation-invariant texture measures for the same test images that were used in recent papers [5,27,28].

This notion of using a subset of all available patterns can be extended further into determining a problem or application specific set of patterns, which provides the best performance. However, there is a danger that the chosen subset becomes too specialized to the training data, and does not generalize properly to unknown testing data.

6 Discussion

The role of local features (primitives, textural elements, textons) and their differences in preattentive texture discrimination has been established in psychophysical studies [1,10]. Our approach fits in well with these findings: simple feature detectors are used to measure perceptually significant textural properties and the segmentation is performed by comparing textural similarities of neighboring regions on the basis of densities of different features.

In our approach, the distributions of texture measures based on local binary patterns and signed gray level differences are used to provide complementary information about the structural and stochastic properties of image texture, making it very efficient for various types of micro- and macrotextures.

The texture primitives detected by LBP could be linked to more complex textural elements using principles of proximity and good continuation, but our excellent results suggest that this is usually not so important. Coarseness, edge orientation and contrast are perceptually very important textural properties [33,12]. The LBP histograms contain information about edge orientation and coarseness (e.g. edges per unit area). LBP does not address the contrast of texture. For this purpose, we can combine LBP with a simple contrast measure C . This measure is invariant with respect to average luminance, which is very important considering, for example, the needs of texture segmentation (see Fig. 5).

LBP is invariant against any monotonic gray scale transformation, which is a necessary property in many applications involving uneven illumination or great within-class variability; see, for example, our results in metal strip inspection [25] and in the classification of tilted 3D textures [32]. The basic LBP method is

sensitive to texture orientation, but this property appears to be very useful in texture segmentation, e.g. for finding accurate boundaries between neighboring regions. A rotation invariant version of LBP was proposed in [16,27], but recently we have significantly improved it by concentrating only on a subset of rotation invariant patterns encoded in LBP and by quantizing the angular space at a finer resolution [21].

The signed gray level difference method is closely related to the widely used co-occurrence method. It is invariant with respect to average luminance, while the co-occurrence approach is not. The method is highly efficient in applications involving stochastic textures, but it also discriminates more structured textures very well. The difference histograms provide information about the coarseness and contrast of texture. For a coarse texture, the difference histogram will peak near zero, but for a finer-grained texture it will spread out more uniformly. The eight-dimensional signed gray level difference operator can be regarded as a generalization of the LBP operator, capturing information about both structural and stochastic aspects of texture.

The importance of the spatial pattern and contrast information in texture discrimination is also supported by our other findings. A joint pair of a rotation-invariant version of LBP and a local variance measure provided the best performance in our experiments with rotation-invariant texture classification [27,21]. In addition, the original LBP measure and a measure based on absolute gray level differences were the best features for two different test image sets in our experiments with multichannel texture description, providing together nearly as good performance as an 'optimal' combination of all 12 measures used in the study [17].

The choice of highly discriminating texture features is very important for success in texture segmentation. The features should easily discriminate various types of textures, and the window size used for computing textural features should be small enough to be useful for small image regions and to provide small error rates at region boundaries. The LBP, LBP:C and signed gray level difference features perform well also for small image regions (e.g., 16 x 16 pixels), unlike many other approaches.

Local texture similarity can be measured very efficiently and reliably with a nonparametric statistical test as in our segmentation procedure. One major advantage is that these methods do not require the specification of a suitable vector-space metric. Similarity is instead defined directly via the respective feature distributions [9].

7 Conclusion

A nonparametric approach to texture analysis using simple spatial operators has been developed. The distributions of texture measures based on local binary patterns and signed gray level differences are used to provide complementary information about the structural and stochastic properties of image texture. Local texture similarity is measured very efficiently and reliably with nonparametric statistical tests. Our methodology resembles some well-known psychophysical models on human texture perception. The approach has provided excellent results in various texture classification and segmentation problems, and leads to computationally simple implementations. Our results should therefore be of great interest from both the theoretical and practical viewpoints when trying to find solutions that could increase the usefulness of texture in real-world applications.

8 Acknowledgements

The financial support provided by the Academy of Finland and the National Technology Agency (Teles) is gratefully acknowledged.

9 Note

Most of the image data used in our papers can be downloaded from <http://www.ee.oulu.fi/research/imag/texture>.

10 References

1. J. Beck, S. Prazdny and A. Rosenfeld, "A theory of textural segmentation," in *Human and Machine Vision*, eds. J. Beck, B. Hope and A. Rosenfeld, Academic, New York, 1983.
2. P. Brodatz, *Textures: A Photographic Album for Artists and Designers*, Dover Publications, New York, 1966.
3. R. Chellappa, R.I. Kashyap and B.S. Manjunath, "Model-based texture segmentation and classification," in *Handbook of Pattern Recognition and Computer Vision, Second Edition*, eds. C.H. Chen, I.F. Pau, P.S.P. Wang, World Scientific, Singapore, 1999, pp. 249-282.
4. L. Van Gool, P. Dewaele and A. Oosterlinck, "Texture analysis anno 1983," *Comput. Vision Graph. Image Process.* 29 (1985) 336-357.

5. G.M. Haley and B.S. Manjunath, "Rotation-invariant texture classification using a complete space-frequency model," *IEEE Trans. Image Process.* 8 (1999) 255-269.
6. R.M. Haralick and L.G. Shapiro, *Computer and Robot Vision, Vol. 1*, Addison-Wesley, 1992.
7. D. Harwood, T. Ojala, M. Pietikäinen, S. Kelman and L. Davis, "Texture classification by center-symmetric autocorrelation, using Kullback discrimination of distributions," *Pattern Recogn. Lett.* 16 (1995) 1-10.
8. D.C. He and L. Wang, "Texture unit, texture spectrum and texture analysis," *IEEE Trans. Geosci. Remote Sensing* 28 (1990) 509-512.
9. T. Hofmann, J. Puzicha and J.M. Buhmann, "Unsupervised texture segmentation in a deterministic annealing framework," *IEEE Trans. Pattern Anal. Mach. Intell.* 20 (1998) 803-818.
10. B. Julesz and J.R. Bergen, "Textons, the fundamental elements in preattentive vision and perception of textures," *The Bell Syst. Techn. J.* 62 (1983) 1619-1645.
11. T. Kohonen, J. Kangas, J. Laaksonen and K. Torkkola, "LVQ_PAK: A program package for the correct application of learning vector quantization algorithms," *Proc. Int. Joint Conf. on Neural Networks*, Baltimore, 1992, pp. 1725-1730.
12. M. Levine, *Vision in Man and Machine*, McGraw-Hill, 1985.
13. B.S. Manjunath and W.Y. Ma, "Texture features for browsing and retrieval of image data," *IEEE Trans. Pattern Anal. Mach. Intell.* 18 (1996) 837-842.
14. T. Mäenpää, T. Ojala, M. Pietikäinen and M. Soriano, "Robust texture classification by subsets of Local Binary Patterns," *Proc. 15th Int. Conf. on Pattern Recognition*, Barcelona, Spain 2000, in press.
15. T. Mäenpää, M. Pietikäinen and T. Ojala, "Texture classification by multi-predicate Local Binary Pattern operators," *Proc. 15th Int. Conf. on Pattern Recognition*, 2000, in press.
16. T. Ojala, Nonparametric Texture Analysis Using Simple Spatial Operators, with Applications in Visual Inspection, *Acta Univ. Ouluensis, C* 105, 1997.
17. T. Ojala and M. Pietikäinen, "Nonparametric multichannel texture description with simple spatial operators," *Proc. 14th Int. Conf. on Pattern Recognition*, Brisbane, Australia, 1998, pp. 1952-1056.
18. T. Ojala and M. Pietikäinen, "Unsupervised texture segmentation using feature distributions," *Pattern Recogn.* 32 (1999) 477-486.
19. T. Ojala, M. Pietikäinen and D. Harwood, "A comparative study of texture measures with classification based on feature distributions," *Pattern Recogn.* 29 (1996) 51-59.
20. T. Ojala, M. Pietikäinen and J. Kyllönen, "Gray level cooccurrence histograms via learning vector quantization," *Proc. 11th Scand. Conf. on Image Analysis*, Kangerlussuaq, Greenland, 1999, pp. 103-108.

21. T. Ojala and M. Pietikäinen and T. Mäenpää, "Gray scale and rotation invariant texture classification with Local Binary Patterns," *Proc. Sixth European Conf. on Computer Vision*, Dublin, Ireland, 2000, in press.
22. T. Ojala, M. Pietikäinen and J. Nisula, "Determining composition of grain mixtures by texture classification based on feature distributions," *Int. J. Pattern Recogn. Artif. Intell.* 10 (1996) 73-82.
23. T. Ojala, K. Valkealahti, E. Oja and M. Pietikäinen, "Texture discrimination with multidimensional distributions of signed gray level differences," *Pattern Recogn. (2000)*, in press.
24. D.K. Panjwani and G. Healey, "Markov random field models for unsupervised segmentation of textured color images," *IEEE Trans. Pattern Anal. Mach. Intell.* 17 (1995) 939-954.
25. M. Pietikäinen, T. Ojala, J. Nisula and J. Heikkinen, "Experiments with two industrial problems using texture classification based on feature distributions," *SPIE 2354 Intelligent Robots and Computer Vision XIII: 3D Vision, Product Inspection and Active Vision*, 1994, pp. 197-204.
26. M. Pietikäinen, T. Ojala and O. Silven, "Approaches to texture-based classification, segmentation and surface inspection," in *Handbook of Pattern Recognition and Computer Vision, Second Edition*, eds. C.H. Chen, L.F. Pau, P.S.P. Wang, World Scientific, Singapore, 1999, pp. 711-736.
27. M. Pietikäinen, T. Ojala and Z. Xu, "Rotation-invariant texture classification using feature distributions," *Pattern Recogn.* 33 (2000), 43-52.
28. R. Porter and N. Canagarajah, "Robust rotation-invariant texture classification: wavelet, Gabor filter and GMRF based schemes," *IEE Proc. Vision Image Signal Process.* 144 (1997) 180-188.
29. J. Puzicha, T. Hofmann and J.M. Buhmann, "Histogram clustering for unsupervised segmentation and image retrieval," *Pattern Recogn. Lett.* 20 (1999) 899-909.
30. T. Randen and J.H. Husøy, "Filtering for texture classification: a comparative study," *IEEE Trans. Pattern Anal. Mach. Intell.* 21 (1999), 291-310. <http://www.ux.his.no/~tranden/>.
31. G. Smith and I. Burns, "Measuring texture classification algorithms," *Pattern Recogn. Lett.* 18 (1997) 1495-1501. <http://www.cssip.elec.uq.edu.au/~guy/meastex/meastex.html>.
32. M. Soriano, T. Ojala and M. Pietikäinen, "Robustness of Local Binary Pattern (LBP) operators to tilt-compensated textures", in *Texture Analysis in Machine Vision*, ed. M. Pietikäinen, World Scientific, Singapore, 2000.
33. H. Tamura, S. Mori and T. Yamawaki, "Textural features corresponding to visual perception," *IEEE Trans. Syst. Man Cybern.* 8 (1978) 460-473.
34. M. Tuceryan and A.K. Jain, "Texture analysis," in *Handbook of Pattern Recognition and Computer Vision, Second Edition*, eds. C.H. Chen, L.F. Pau, P.S.P. Wang, World Scientific, Singapore, 1999, pp. 207-248.

35. K. Valkealahti and E. Oja, "Reduced multidimensional cooccurrence histograms in texture classification," *IEEE Trans. Pattern Anal. Mach. Intell.* 20 (1998) 90-94.